Cache-Aided Coded Multicast for Correlated Sources

P. Hassanzadeh, A. Tulino, J. Llorca, E. Erkip

arXiv:1609.05831v1 [cs.IT] 19 Sep 2016

***Abstract*—The combination of edge caching and coded multicasting is a promising approach to improve the efficiency of content delivery over cache-aided networks. The global caching gain resulting from content overlap distributed across the network in current solutions is limited due to the increasingly personalized nature of the content consumed by users. In this paper, the cache-aided coded multicast problem is generalized to account for the correlation among the network content by formulating a source compression problem with distributed side information. A correlation-aware achievable scheme is proposed and an upper bound on its performance is derived. It is shown that considerable load reductions can be achieved, compared to state of the art correlation-unaware schemes, when caching and delivery phases specifically account for the correlation among the content files.**

I. Introduction

Proper distribution of popular content across the network caches is emerging as one of the promising approaches to address the exponentially growing traffic in current wireless networks. Recent studies [1]–[6] have shown that, in a cache-aided network, exploiting globally cached information in order to multicast coded messages that are useful to a large number of receivers exhibits overall network throughput that is proportional to the aggregate cache capacity. The fundamental rate-memory trade-off in a broadcast caching network has been characterized in [1]–[6]. While these initial results are promising, these studies treat the network content as independent pieces of information, and do not account for the additional potential gains arising from further compression of correlated content distributed across the network. , In this paper, we investigate how the correlations among the library content can be explored in order to further improve the performance in cache-aided networks. ,We consider a network setup similar to [1]–[6], but assume that the files in the library are correlated. Such correlations are especially relevant among content files of the same category, such as episodes of a TV show or same-sport recordings, which, even if personalized, may share common backgrounds and scene objects.

As in existing literature on cache-aided networks, we assume that the network operates in two phases: a caching (or placement) phase taking place at network setup followed by a delivery phase where the network is used repeatedly in order to satisfy receiver demands. The design of the caching and delivery phases forms what is referred to as a *caching scheme*. During the caching phase, caches are filled with content from the library according to a properly designed *caching distribution*. During the delivery phase, the sender compresses the set of requested files into a multicast codeword by computing an *index code* while exploring all correlations that exist among the requested and cached content. In [7], the rate-memory region for a correlated library was characterized for lossy reconstruction in a scenario with only two receivers and a single cache. In this paper, we consider correlation-aware lossless reconstruction in a more general setting with multiple caches, and with a nonuniform demand distribution.

We propose an achievable correlation-aware scheme, named Correlation-Aware RAP caching and Coded Multicast delivery (CA-RAP/CM), in which receivers store content pieces based on their popularity as well as on their correlation with the rest of the file library during the caching phase, and receive compressed versions of the requested files according to the information distributed across the network and their joint statistics during the delivery phase. The scheme consists of: 1) exploiting file correlations to store the *more relevant bits* during the caching phase such that expected delivery rate is reduced, and 2) optimally designing the coded multicast codeword based on the joint statistics of the library files and the aggregate cache content during the delivery phase. Additional refinements are transmitted, when needed, in order to ensure lossless reconstruction of the requested files at each receiver. Given the exponential complexity of CA-RAP/CM, we then provide an algorithm which approximates CA-RAP/CM in polynomial time, and we derive an upper bound on the achievable expected rate. We numerically compare the rates achieved by the proposed correlation-aware scheme with existing correlation-unaware schemes, and our results confirm the additional gains achievable, especially for small memory size.

The paper is organized as follows. The problem formulation is presented in Sec. II. In Sec. III we describe CA-RAP/CM and its polynomial-time approximation, and we quantify the associated rate-memory trade-off. We provide numerical simulations and concluding remarks in Secs. IV and V, respectively.

II. Network Model and Problem Formulation

Notation: For ease of exposition, we use $\{A_i\}$ to denote the set of elements $\{A_i : i \in \mathcal{I}\}$, with $\mathcal{I}$ being the domain of index i, and we define $[n] \triangleq \{1, \dots, n\}$. $\mathbb{F}_2^*$ denotes the set of finite length binary sequences.

Similar to previous work [3], [4], we consider a broadcast caching network composed of one sender (e.g., base station) with access to a file library composed of m files, each of entropy F bits. Without loss of generality, we assume that each file $f \in [m]$ is represented by a vector of i.i.d. random binary symbols of length F, $\mathsf{W}_f \in \mathbb{F}_2^F$, and

P. Hassanzadeh and E. Erkip are with the ECE Department of New York University, Brooklyn, NY. Email: {ph990, elza}@nyu.edu

J. Llorca and A. Tulino are with Bell Labs, Nokia, Holmdel, NJ, USA. Email: {jaime.llorca, a.tulino}@nokia.com

A. Tulino is with the DIETI, University of Naples Federico II, Italy. Email: {antoniamaria.tulino}@unina.it

hence $H(\mathsf{W}_f) = F$ for all $f \in [m]$.[1] We assume that the symbols belonging to different files can be correlated, i.e., $H(\mathsf{W}_1, \ldots, \mathsf{W}_m) \leq mF$, and we denote their joint distribution by $P_{\mathcal{W}}$. The sender communicates with n receivers (e.g., access points or user devices) $\mathcal{U} = \{1, \ldots, n\}$ through a shared error-free multicast link. Each receiver has a cache of size MF bits and requests files in an independent and identically distributed (i.i.d.) manner according to a demand distribution $\mathbf{q} = (q_1, \ldots, q_m)$, where q_f denotes the probability of requesting file $f \in [m]$. A caching scheme for this network consists of:

- **Cache Encoder:** Given a library realization $\{W_f\}$, the cache encoder at the sender computes the cache content $Z_u(\{W_f\})$ at receiver $u \in \mathcal{U}$, using the set of functions $\{Z_u : \mathbb{F}_2^{mF} \to \mathbb{F}_2^{MF} : u \in \mathcal{U}\}$. The encoder designs the cache configuration $\{Z_u\}$ jointly across receivers, taking into account global system knowledge such as the number of receivers and their cache sizes, the number of files, their aggregate popularity, and their joint distribution $P_{\mathcal{W}}$.
- **Multicast Encoder:** After the caches are populated, the network is repeatedly used for different demand realizations. At each use of the network, the receivers place a random demand vector $\mathbf{f} \in [m]^n$ at the sender, according to $\mathbf{q}$. The demand realization is denoted by $\mathbf{f} = (f_1, \ldots, f_n)$. The multicast encoder at the sender computes codeword $X(\mathbf{f}, \{W_f\}, \{Z_u\})$, which will be transmitted over the shared link, using the function $X : [m]^n \times \mathbb{F}_2^{mF} \times \mathbb{F}_2^{nMF} \to \mathbb{F}_2^*$. In this work, we consider a fixed-to-variable almost-lossless framework.
- **Multicast Decoders:** Receivers recover their requested files using their cached content and the received multicast codeword. More specifically, receiver u recovers W_{f_u} using its decoding function $\zeta_u : [m]^n \times \mathbb{F}_2^* \times \mathbb{F}_2^{MF} \to \mathbb{F}_2^F$, as $\widehat{W}_{f_u} = \zeta_u(\mathbf{f}, X, Z_u)$.

The worst-case (over the file library) probability of error of the corresponding caching scheme is defined as

$$P_e^{(F)} = \sup_{\{W_f : f \in [m]\}} \mathbb{P}\left(\widehat{W}_{f_u} \neq W_{f_u}\right).$$

In line with previous work [3], [4], the (average) rate of the overall caching scheme is defined as

$$R^{(F)} = \sup_{\{W_f : f \in [m]\}} \frac{\mathbb{E}[J(X)]}{F}. \tag{1}$$

where $J(X)$ denotes the length (in bits) of the multicast codeword X.

Definition 1. *A rate R is* achievable *if there exists a sequence of caching schemes for increasing file size F such that $\lim_{F\to\infty} P_e^{(F)} = 0$, and $\limsup_{F\to\infty} R^{(F)} \leq R$.*

The goal of this paper is to design a caching scheme that results in a small achievable rate R.

[1] The results derived in this paper can be easily extended to the case in which the symbols within each file are not necessarily binary and independent.

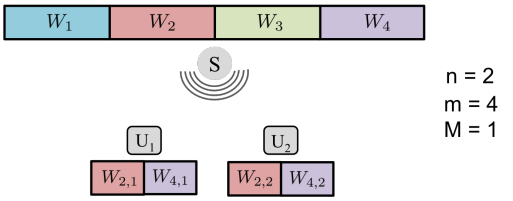

Fig. 1: Network setup and cache configuration of Example 1.

III. Correlation-Aware RAP Caching and Coded Multicast Delivery (CA-RAP/CM)

In this section, we introduce CA-RAP/CM, a correlation-aware caching scheme, which is an extension of a RAndom Popularity-based (RAP) caching policy followed by a Chromatic-number Index Coding (CIC) delivery policy [3], [4]. In CA-RAP/CM, both the cache encoder and the multicast encoder are designed according to the joint distribution $P_{\mathcal{W}}$, in order to exploit the correlation among the library files. First, consider the following motivating example.

Example 1. *Consider a file library with $m = 4$ uniformly popular files $\{W_1, W_2, W_3, W_4\}$ each with entropy F bits. We assume that the pairs $\{W_1, W_2\}$ and $\{W_3, W_4\}$ are independent, while correlations exist between W_1 and W_2, and between W_3 and W_4. Specifically, $H(\mathsf{W}_1|\mathsf{W}_2) = \mathsf{H}(\mathsf{W}_2|\mathsf{W}_1) = \mathsf{F}/4$ and $H(\mathsf{W}_3|\mathsf{W}_4) = \mathsf{H}(\mathsf{W}_4|\mathsf{W}_3) = \mathsf{F}/4$. The sender is connected to $n = 2$ receivers $\{u_1, u_2\}$ with cache size $M = 1$. While a correlation-unaware scheme (e.g., [2], [3]) would first compress the files separately and then cache $1/4^{th}$ of each file at each receiver, existing file correlations can be exploited to cache the more relevant bits. For example, one could split files W_2 and W_4 into two parts $\{W_{2,1}, W_{2,2}\}$ and $\{W_{4,1}, W_{4,2}\}$, each with entropy $F/2$, and cache $\{W_{2,1}, W_{4,1}\}$ at u_1 and $\{W_{2,2}, W_{4,2}\}$ at u_2, as shown in Fig. 1. During the delivery phase, consider the worst case demand, e.g., $\mathbf{f} = (W_3, W_1)$, the sender first multicasts the XOR of the compressed parts $W_{2,1}$ and $W_{4,2}$. Refinement segments, with refinement rates $H(\mathsf{W}_3|\mathsf{W}_4)$ and $H(\mathsf{W}_1|\mathsf{W}_2)$ can then be transmitted to enable lossless reconstruction, resulting in a total rate $R = 1$. Note that a correlation-unaware scheme would need a total rate $R = 1.25$ regardless of the demand realization [2], [3].*

The main components of the CA-RAP/CM scheme are: i) a *Correlation-Aware Random Popularity Cache Encoder* (CA-RAP) and ii) a *Correlation-Aware Coded Multicast Encoder* (CA-CM).

A. Correlation-Aware Random Popularity Cache Encoder

The CA-RAP cache encoder is a correlation-aware random fractional cache encoder whose key differentiation from the cache encoder RAP, introduced in [3], [4], is to choose the fractions of files to be cached according to both their popularity as well as their correlation with the rest of the library. Similar to the cache encoder RAP, each file is partitioned into B equal-size packets, with packet $b \in [B]$ of file $f \in [m]$ denoted by $W_{f,b}$. The cache content at each receiver is selected according to a *caching distribution*, $\mathbf{p} = (p_1, \ldots, p_m)$

with $0 \leq p_f \leq 1/M \ \forall f \in [m]$ and $\sum_{f=1}^{m} p_f = 1$, which is optimized to minimize the rate of the corresponding index coding delivery scheme. For a given caching distribution $\mathbf{p}$, each receiver caches a subset of $p_f MB$ distinct packets from each file $f \in [m]$, independently at random. We denote by $\mathbf{C} = \{\mathbf{C}_1, \ldots, \mathbf{C}_n\}$ the packet-level cache configuration, where $\mathbf{C}_u$ denotes the set of file-packet index pairs, (f, b), $f \in [m]$, $b \in [B]$, cached at receiver u. In Example 1, $B = 2$, the caching distribution corresponds to $p_{W_2} = p_{W_4} = 1/2$, $p_{W_1} = p_{W_3} = 0$, and the packet-level cache configuration is $\mathbf{C} = \{\{(2,1),(4,1)\},\{(2,2),(4,2)\}\}$. While the caching distribution of a correlation-unaware scheme prioritizes the caching of packets according to the aggregate popularity distribution (see [3], [4]), the CA-RAP caching distribution accounts for both the aggregate popularity and the correlation of each file with the rest of the library when determining the amount of packets to be cached from each file.

B. Correlation-Aware Coded Multicast Encoder

For a given demand realization $\mathbf{f}$, the packet-level demand realization is denoted by $\mathbf{Q} = [\mathbf{Q}_1, \ldots, \mathbf{Q}_n]$, where $\mathbf{Q}_u$ denotes the file-packet index pairs (f, b) associated with the packets of file W_{f_u} requested, but not cached, by receiver u.

The CA-CM encoder capitalizes on the additional coded multicast opportunities that arise from incorporating cached packets that are, not only equal to, but also correlated with the requested packets into the multicast codeword. The CA-CM encoder operates by constructing a *clustered conflict graph*, and computing a linear index code from a valid[2] coloring of the conflict graph, as described in the following.

1) Correlation-Aware Clustering: For each requested packet $W_{f,b}$, $(f, b) \in \mathbf{Q}$, the correlation-aware clustering procedure computes a *δ-ensemble* $\mathcal{G}_{f,b}$, where $\mathcal{G}_{f,b}$ is the union of $W_{f,b}$ and the subset of all cached and requested packets that are δ-correlated with $W_{f,b}$, as per the following definition.[3]

Definition 2. *(δ-**Correlated Packets**) For a given threshold $\delta \leq 1$, packet $\mathsf{W}_{f,b}$ is δ-correlated with packet $\mathsf{W}_{f',b'}$ if $H(\mathsf{W}_{f,b}, \mathsf{W}_{f',b'}) \leq (1+\delta)F$ bits, for all $f, f' \in [m]$ and $b, b' \in [B]$.*

This classification (clustering) is the first step for constructing the clustered conflict graph.

2) Correlation-Aware Cluster Coloring: The clustered conflict graph $\mathcal{H}_{\mathbf{C},\mathbf{Q}} = (\mathcal{V}, \mathcal{E})$ is constructed as follows:

- Vertex set $\mathcal{V}$: The vertex (node) set $\mathcal{V} = \widehat{\mathcal{V}} \cup \widetilde{\mathcal{V}}$ is composed of root nodes $\widehat{\mathcal{V}}$ and virtual nodes $\widetilde{\mathcal{V}}$.
  - Root Nodes: There is a root node $\hat{v} \in \widehat{\mathcal{V}}$ for each packet $W_{f,b}$, requested by each receiver, uniquely identified by the pair $\{\rho(\hat{v}), \mu(\hat{v})\}$, with $\rho(\hat{v})$ denoting the packet identity, i.e., file-packet index pair (f, b), and $\mu(\hat{v})$ denoting the receiver requesting it.
  - Virtual Nodes: For each root node $\hat{v} \in \widehat{\mathcal{V}}$, all the packets in the δ-ensemble $\mathcal{G}_{\rho(\hat{v})}$ other than $\rho(\hat{v})$ are represented as virtual nodes in $\widetilde{\mathcal{V}}$. We identify virtual node $\tilde{v} \in \widetilde{\mathcal{V}}$, having $\hat{v}$ as a root note, with the triplet $\{\rho(\tilde{v}), \mu(\hat{v}), r(\tilde{v})\}$, where $\rho(\tilde{v})$ indicates the packet identity associated with vertex $\tilde{v}$, $\mu(\hat{v})$ indicated the receiver requesting $\rho(\hat{v})$, and $r(\tilde{v}) = \hat{v}$ is the root of the δ-ensemble that $\tilde{v}$ belongs to.

  We denote by $\mathcal{K}_{\hat{v}} \subseteq \mathcal{V}$ the set of vertices containing root node $\hat{v} \in \widehat{\mathcal{V}}$ and the virtual nodes corresponding to the packets in its δ-ensemble $\mathcal{G}_{\rho(\hat{v})}$, and we refer to $\mathcal{K}_{\hat{v}}$ as the *cluster* of root node $\hat{v}$.
- Edge set $\mathcal{E}$: For any pair of vertices $v_1, v_2 \in \mathcal{V}$, there is an edge between v_1 and v_2 in $\mathcal{E}$ if both v_1 and v_2 are in the same cluster, or if the two following conditions are jointly satisfied 1) $\rho(v_1) \neq \rho(v_2)$, 2) packet $\rho(v_1) \notin \mathbf{C}_{\mu(v_2)}$ or packet $\rho(v_2) \notin \mathbf{C}_{\mu(v_1)}$.

  Given a valid coloring of $\mathcal{H}_{\mathbf{C},\mathbf{Q}}$, a valid *cluster coloring* of $\mathcal{H}_{\mathbf{C},\mathbf{Q}}$ consists of assigning to each cluster $\mathcal{K}_{\hat{v}}$, $\forall \hat{v} \in \widehat{\mathcal{V}}$, one of the colors assigned to the vertices inside that cluster. For each color in the cluster coloring, only the packets with the same color as the color assigned to their corresponding cluster, are XORed together and multicasted to the users. Using its cached information and the received XORed packets, each receiver is able to reconstruct a (possibly) distorted version of its requested packet, due to the potential reception of a packet that is δ-correlated with its requested one. The encoder transmits refinement segments, when needed, to enable lossless reconstruction of the demand at each receiver. The coded multicast codeword results from concatenating: 1) for each color in the cluster coloring, the XOR of the packets with the same color, and, 2) for each receiver, if needed, the refinement segment. The CA-CM encoder selects the valid cluster coloring corresponding to the shortest coded multicast codeword.

Remark 1. *Note that if correlation is not considered or non-existent, the clustered conflict graph is equivalent to the conventional index coding conflict graph [3], [4], i.e., the subgraph of $\mathcal{H}_{\mathbf{C},\mathbf{Q}}$ resulting from considering only the root nodes $\widehat{\mathcal{V}}$.*

Remark 2. *It is important to note that the number of colors in the cluster coloring chosen by CA-CM is always smaller than or equal to the chromatic number[4] of the conventional index coding conflict graph. Such reduction in the number of colors is obtained by considering correlated packets that are cached in the network, which possibly results in less conflict edges and provides more options for coloring each cluster. Intuitively, CA-CM allows for the requested packets that had to be transmitted by themselves otherwise, to be represented by correlated packets that can be XORED together with other packets in the multicast codeword.*

[2]A valid coloring of a graph is an assignment of colors to the vertices of the graph such that no two adjacent vertices are assigned the same color.

[3]In practice, the designer shall determine the level at which to compute and exploit correlations between files based on performance-complexity tradeoffs. For example, in video applications, a packet may represent a video block or frame.

[4]The chromatic number of a graph is the minimum number of colors over all valid colorings of the graph.

C. Greedy Cluster Coloring (GClC)

Given that graph coloring, and by extension cluster coloring, is NP-Hard, in this section we propose *Greedy Cluster Coloring (GClC)*, a polynomial-time approximation to the cluster coloring problem. GClC extends the existing Greedy Constraint Coloring (GCC) scheme [4], to account for file correlation in cache-aided networks and consists of a combination of two coloring schemes, such that the scheme resulting in the lower number of colors (i.e., shorter multicast codeword) is chosen. Uncoded refinement segments are transmitted to ensure lossless reconstruction of the demand.

Algorithm 1 GClC_1 - Greedy Cluster Coloring 1

1: **while** $\widehat{\mathcal{V}} \neq \emptyset$ **do**
2: Pick any root node $\hat{v} \in \widehat{\mathcal{V}}$
3: Sort $\mathcal{K}_{\hat{v}}$ in decreasing order of receiver label size[5], such that for $v_t, v_{t+1} \in \mathcal{K}_{\hat{v}}$, $|\{\mu(v_t), \eta(v_t)\}| \geq |\{\mu(v_{t+1}), \eta(v_{t+1})\}|$ where v_t denotes the t^{th} vertex in the ordered sequence.
4: $t = 1$
5: **while** $t \leq |\mathcal{K}_{\hat{v}}|$ **do**
6: Take $v_t \in \mathcal{K}_{\hat{v}}$; Let $I_{v_t} = \{v_t\}$
7: **for all** $v \in \mathcal{V} \backslash \Big\{\mathcal{K}_{\hat{v}} \cup I_{v_t}\Big\}$ **do**
8: **if** $\Big\{$ There is no edge between v and I_{v_t} $\Big\} \bigwedge$ $\Big\{\{\mu(v), \eta(v)\} = \{\mu(v_t), \eta(v_t)\}\Big\}$ **then**
9: $I_{v_t} = I_{v_t} \cup \{v\}$
10: **end if**
11: **end for**
12: ${v_t}^* = \underset{\{v_\tau : \tau = 1, \ldots, t\}}{\arg\max} |I_{v_\tau}|$
13: **if** $|I_{v_t^*}| \geq |\{\mu(v_t), \eta(v_t)\}|$ or $t = |\mathcal{K}_{\hat{v}}|$ **then**
14: $\mathcal{I} = I_{v_t^*}$
15: $t = |\mathcal{K}_{\hat{v}}| + 1$
16: **else**
17: $t = t + 1$
18: **end if**
19: **end while**
20: Color all vertices in $\mathcal{I}$ with an unused color.
21: $\mathcal{J} = \{r(v) : v \in \mathcal{I}\}$
22: $\widehat{\mathcal{J}} = \Big\{\hat{v} \in \widehat{\mathcal{V}} : \exists v \in \mathcal{I}, \Big\{\mu(\hat{v}) = \mu(v)\Big\} \bigwedge \Big\{\rho(v) \in \mathcal{G}_{\rho(\hat{v})}\Big\}\Big\}$
23: $\widehat{\mathcal{V}} \leftarrow \widehat{\mathcal{V}} \backslash \mathcal{J} \cup \widehat{\mathcal{J}}$, $\mathcal{V} \leftarrow \mathcal{V} \backslash \bigcup_{\hat{v} \in \mathcal{J} \cup \widehat{\mathcal{J}}} \mathcal{K}_{\hat{v}}$
24: **end while**

In GClC, we assume that any vertex (root node or virtual node) $v \in \mathcal{V}$ is identified by the triplet $\{\rho(v), \mu(v), r(v)\}$, which is uniquely specified by the packet identity associated with v and by the cluster to which v belongs. Specifically, given a vertex $v \in \mathcal{K}_{\hat{v}}$, then $\rho(v)$ indicates the packet identity associated with vertex v, while $\mu(v) = \mu(\hat{v})$ and $r(v) = \hat{v}$. Further define $\eta(v) \triangleq \{u \in \mathcal{U} : \rho(v) \in \mathbf{C}_u\}$ for any $v \in \mathcal{V}$. The unordered set of receivers $\{\mu(v), \eta(v)\}$, corresponding to the set of receivers either requesting or caching packet $\rho(v)$, is referred to as the *receiver label* of vertex v.

Algorithm GClC_1 starts from a root node $\hat{v} \in \widehat{\mathcal{V}}$ among those not yet selected, and searches for the node $v_t \in \mathcal{K}_{\hat{v}}$ which forms the largest independent set $\mathcal{I}$ with all the vertices in $\mathcal{V}$ having its same receiver label.[6] Next, vertices in set $\mathcal{I}$ are assigned the same color (see lines 20-23). Algorithm GClC_2 is based on a correlation-aware extension of GCC_2 in [4], and corresponds to a generalized uncoded (naive) multicast: For each root node $\hat{v} \in \widehat{\mathcal{V}}$, whose cluster has not yet been colored, only the vertex $v_t \in \mathcal{K}_{\hat{v}}$ whom is found among the nodes of more clusters, i.e., correlated with a larger number of requested packets, is colored and its color is assigned to $\mathcal{K}_{\hat{v}}$ and all clusters containing v_t. For both GClC_1 and GClC_2, when the graph coloring algorithm terminates, only a subset of the graph vertices, $\mathcal{V}$, are colored such that only one vertex from each cluster in the graph is colored. This is equivalent to identifying a valid cluster coloring where each cluster is assigned the color of its colored vertex. Between GClC_1 and GClC_2, the cluster coloring resulting in the lower number of colors is chosen. For each color assigned during the coloring, the packets with the same color are XORed together, and multicasted.

D. Performance of CA-RAP/CM

In this section, we provide an upper bound of the rate achieved with CA-RAP/CM. For a given δ, we define the *match matrix* $\mathbf{G}$ as the matrix whose element $\mathbf{G}_{f'f}$ $(f, f') \in [m]^2$ is the largest value such that for each packet $W_{f,b}$ of file f, there are at least $\mathbf{G}_{f'f}$ packets of file f' that are δ-correlated with $W_{f,b}$, and are distinct from the packets correlated with packet $W_{f,b'}$, $\forall\, b' \in [B]$.

Theorem 1. *Consider a broadcast caching network with n receivers, cache capacity M, demand distribution $\mathbf{q}$, a caching distribution $\mathbf{p}$, library size m, correlation parameter δ, and match matrix $\mathbf{G}$. Then, the achievable expected rate of CA-RAP/CM, $R(\delta, \mathbf{p})$, is upper bounded, as $F \to \infty$, with high probability as*

$$R(\delta, \mathbf{p}) \leq \min\left\{\psi(\delta, \mathbf{p}) + \Delta R(\delta, \mathbf{p}), \bar{m}\right\}, \tag{2}$$

where

$$\psi(\delta, \mathbf{p}) \triangleq \sum_{\ell=1}^{n} \binom{n}{\ell} \sum_{f=1}^{m} \rho_{\ell,f} \lambda_{\ell,f},$$

$$\Delta R(\delta, \mathbf{p}) \leq \sum_{\ell=1}^{n} \ell \binom{n}{\ell} \sum_{f=1}^{m} \rho^*_{\ell,f} \lambda^*_{\ell,f} \delta + n \sum_{f=1}^{m} q_f (1 - p_f M) \Big(1 - \prod_{f'=1}^{m} (1 - p_{f'} M)^{(\mathbf{G} - \mathbf{I})_{f'f}}\Big) \delta,$$

$$\bar{m} \triangleq \sum_{f=1}^{m} (1 - (1 - q_f)^n),$$

with

$$\lambda_{\ell,f} \triangleq \Big(\prod_{f'=1}^{m} (1 - p_{f'} M)^{(n-\ell+1)\mathbf{G}_{f'f}}\Big) \Big(1 - \prod_{f'=1}^{m} \Big(1 - (p_{f'} M)^{\ell-1}\Big)^{\mathbf{G}_{f'f}}\Big),$$

[5] Packets with equal label size are ordered such that $H(\mathsf{W}_{\rho(\hat{v})} | \mathsf{W}_{\rho(v_t)}) \leq H(\mathsf{W}_{\rho(\hat{v})} | \mathsf{W}_{\rho(v_{t+1})})$

[6] An independent set is a set of vertices in a graph, no two of which are adjacent.

$$\rho_{\ell,f} \triangleq \mathbb{P}\Big\{ f = \arg\max_{d\in\mathcal{D}} \lambda_{\ell,d} \Big\},$$

$$\lambda^*_{\ell,f} \triangleq \Big(\prod_{f'=1}^{m} (1-p_{f'}M)^{(n-\ell+1)\mathbf{G}_{f'f}} \Big) \Big(1-(p_fM)^{\ell-1}\Big)\Big(1-\prod_{f'=1}^{m}\Big(1-(p_{f'}M)^{\ell-1}\Big)^{(\mathbf{G}-\mathbf{I})_{f'f}}\Big),$$

$$\rho^*_{\ell,f} \triangleq \mathbb{P}\Big\{ f = \arg\max_{d\in\mathcal{D}} \lambda^*_{\ell,d} \Big\},$$

$\mathcal{D}$ denoting a random set of ℓ elements selected in an i.i.d. manner from $[m]$, and $\mathbf{I}$ denoting the identity matrix.

Proof: For proof of Theorem 1 see [8]. ∎

The CA-RAP caching distribution is computed as the minimizer of the corresponding rate upper bound, $\mathbf{p}^* = \arg\min_{\mathbf{p}} R(\delta, \mathbf{p})$, resulting in the optimal CA-RAP/CM rate $R(\delta, \mathbf{p}^*)$. The resulting distribution $\mathbf{p}^*$ may not have an analytically tractable expression in general, but can be numerically optimized for the specific library realization. We remark that the rate upper bound is derived for a given correlation parameter δ, whose value can also be optimized to minimize the achievable expected rate $R(\delta, \mathbf{p})$.

IV. Simulations and Discussions

We numerically compare the performance of the polynomial-time CA-RAP/CM scheme given in Sec. III-D, with existing correlation-unaware schemes: Local Caching with Unicast Delivery (LC/U), Local Caching with Naive Multicast Delivery (LC/NM), and RAP Caching with Coded Multicast Delivery (RAP/CM) [4]. LC/U consists of the conventional LFU caching policy, in which the M most popular files are cached at each receiver, followed by unicast delivery. LC/NM combines LFU caching with naive multicast delivery, where a common uncoded stream of data packets is simultaneously received and decoded by multiple receivers. RAP/CM is the combination of RAP caching and coded multicasting as presented in [4], and is equivalent to CA-RAP/CM when $\delta = 0$.

We consider a broadcast caching network with $n = 10$ receivers and $m = 100$ files requested according to a Zipf distribution with parameter α given by $q_f = f^{-\alpha}/\sum_{j=1}^{m} j^{-\alpha}$, $\forall f \in [m]$. Under demand distribution $\mathbf{q}$, we assume a match matrix $\mathbf{G}$ such that, for each $(f, f') \in [m]^2$, $\mathbf{G}_{f,f'} = B_{\text{match}}$. Fig. 2 displays the rates for $\alpha = 0.8$, $\delta = 0.2$, and $mB_{\text{match}} = 4$, as M varies from 0 to 100. We observe that the correlation-aware scheme is able to achieve rate reductions that go well beyond the state of the art correlation-unaware counterpart. Specifically, CA-RAP/CM achieves a $2.7\times$ reduction in the expected rate compared to LC/U and a $2.4\times$ reduction compared to RAP/CM for $M = 10$.

V. Conclusion

In this work, we have shown how exploiting the correlation among the library files can result in more efficient content delivery over cache-aided networks. We proposed a correlation-aware caching scheme in which receivers store content pieces based on their popularity as well as on their correlation with the rest of the file library in the caching phase, and receive compressed versions of the requested files according to the information distributed across the network and their joint statistics during the delivery phase. The proposed scheme is shown to significantly outperform state of the art approaches that treat library files as mutually independent.

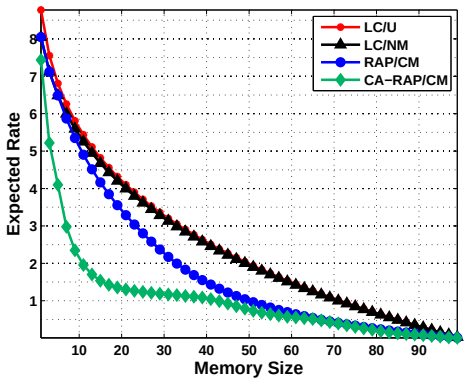

Fig. 2: Performance of CA-RAP/CM in a network with Zipf parameter $\alpha = 0.8$, $n = 10$, $m = 100$, $\delta = 0.2$, and $mB_{\text{match}} = 4$.

Ongoing and future work entail investigating the effect of relevant system parameters that characterize the correlated library, such as δ and $\mathbf{G}$, on the achievable expected rate, as well as providing order-optimality results. An alternative correlation-aware scheme is proposed in [9] where the more relevant content from the jointly compressed library is stored during the caching phase, and the transmissions during the delivery phase ensure perfect reconstruction of the requested content.

References

[1] M. A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *IEEE Transactions on Information Theory*, vol. 60, no. 5, pp. 2856–2867, 2014.

[2] M. Maddah-Ali and U. Niesen, "Decentralized coded caching attains order-optimal memory-rate tradeoff," *IEEE/ACM Transactions on Networking*, no. 99, pp. 1029–1040, 2014.

[3] M. Ji, A. Tulino, J. Llorca, and G. Caire, "On the average performance of caching and coded multicasting with random demands," in *Proc. IEEE International Symposium on Wireless Communications Systems (ISWCS)*, 2014, pp. 922–926.

[4] ——, "Order-optimal rate of caching and coded multicasting with random demands," *arXiv:1502.03124*, 2015.

[5] ——, "Caching-aided coded multicasting with multiple random requests," in *Proc. IEEE Information Theory Workshop (ITW)*, 2015.

[6] M. Ji, K. Shanmugam, G. Vettigli, J. Llorca, A. Tulino, and G. Caire, "An efficient multiple-groupcast coded multicasting scheme for finite fractional caching," in *Proc. IEEE International Conference on Communications (ICC)*, 2015.

[7] R. Timo, S. S. Bidokthi, M. Wigger, and B. Geiger, "A rate-distortion approach to caching," *preprint http://roytimo. wordpress. com/pub*, 2015.

[8] P. Hassanzadeh, A. Tulino, J. Llorca, and E. Erkip, "Rate-memory trade-off for caching and delivery of correlated sources," *preprint http://engineering.nyu.edu/eelab/people/Parisa*.

[9] ——, "Correlation-aware distributed caching and coded delivery," in *Proc. IEEE Information Theory Workshop (ITW)*, 2016.